

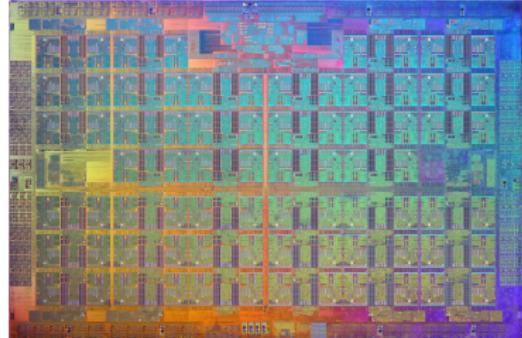
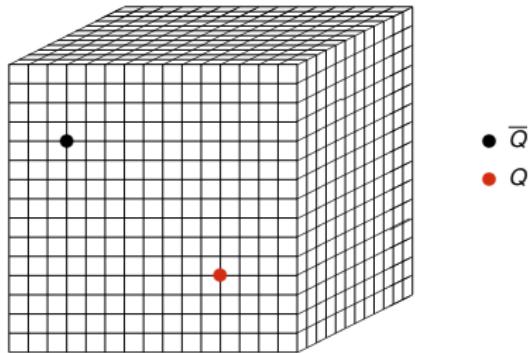


# Lattice QCD on modern CPU architectures

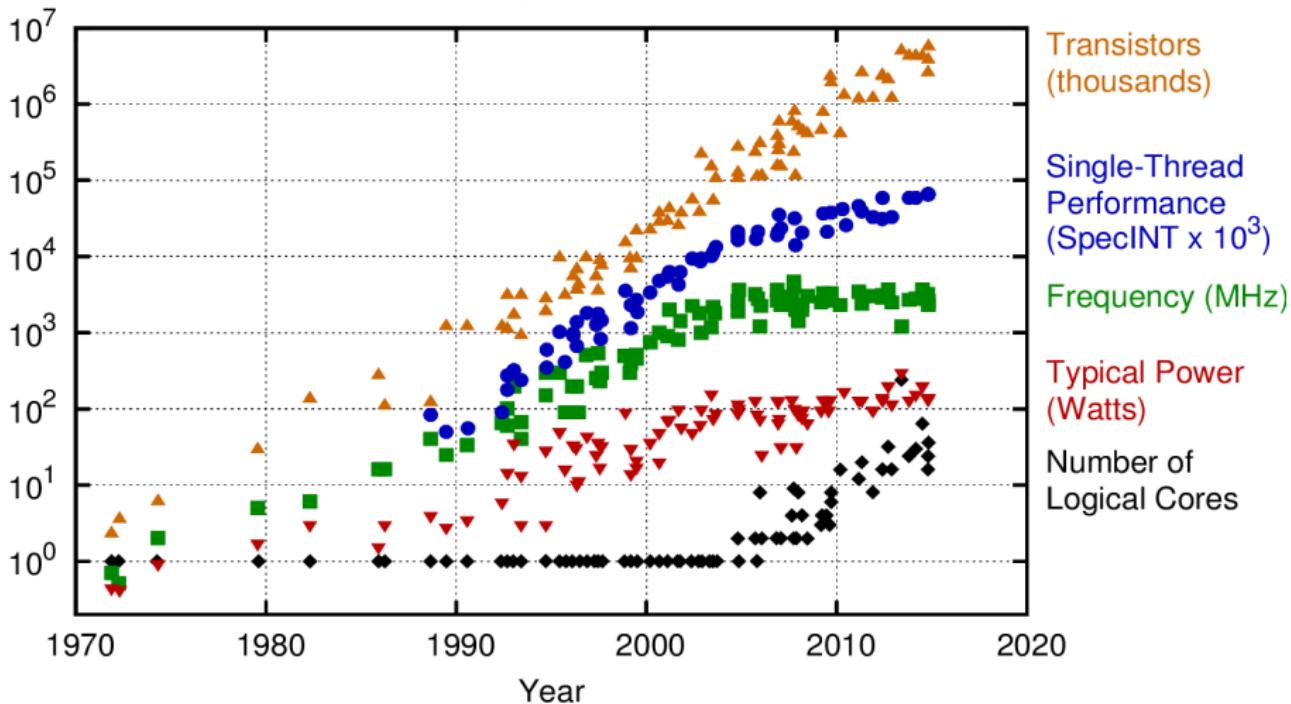
March 10, 2016 | Patrick Steinbrecher

# Outline

- The last decade of x86
- Dslash kernel
- The performance gap:
  - plain C vs. optimized code



## 40 Years of Microprocessor Trend Data



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten  
New plot and data collected for 2010-2015 by K. Rupp

# SIMD: Single Instruction, Multiple Data

$x_0$	$x_1$	$x_2$	$x_3$
-------	-------	-------	-------

+

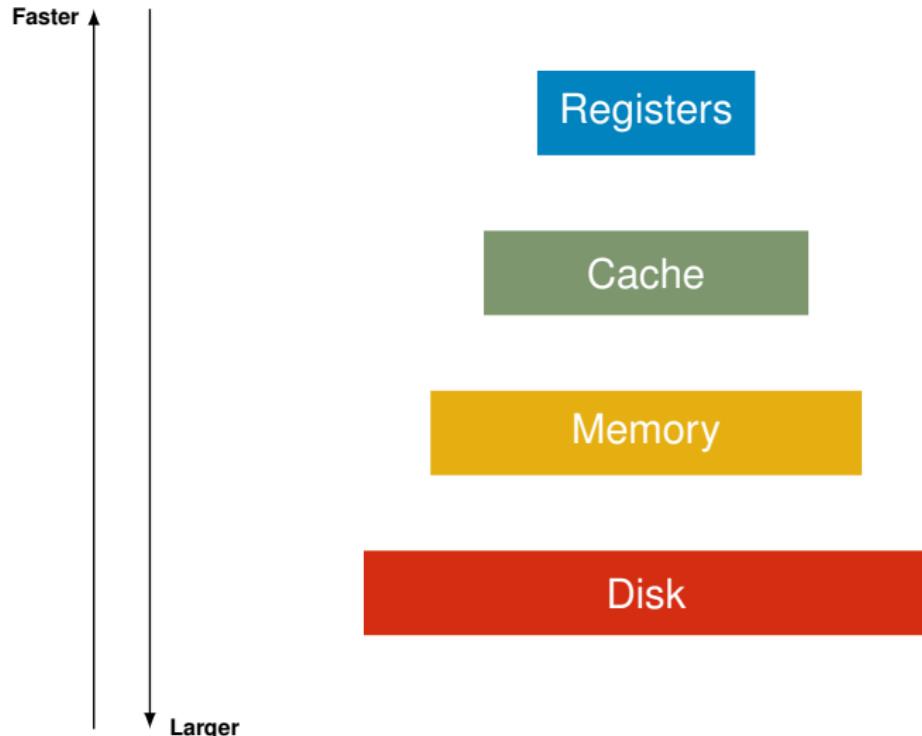
$y_0$	$y_1$	$y_2$	$y_3$
-------	-------	-------	-------

=

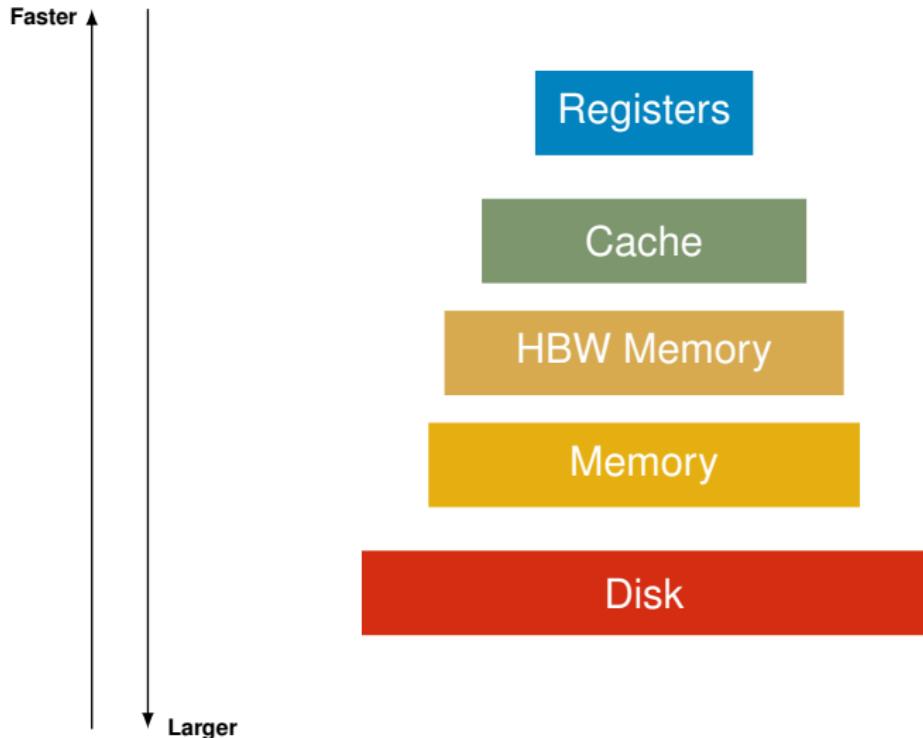
$x_0 + y_0$	$x_1 + y_1$	$x_2 + y_2$	$x_3 + y_3$
-------------	-------------	-------------	-------------

- SISD processor
  - 4 instructions
- SIMD processor
  - 1 instruction

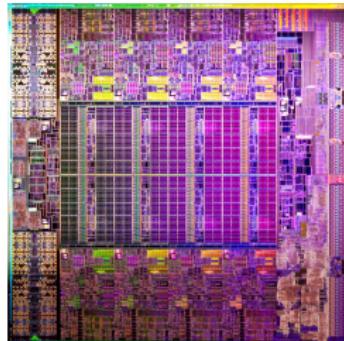
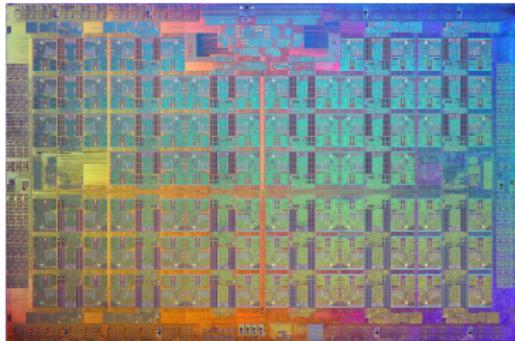
# CPU hierarchy



# CPU hierarchy

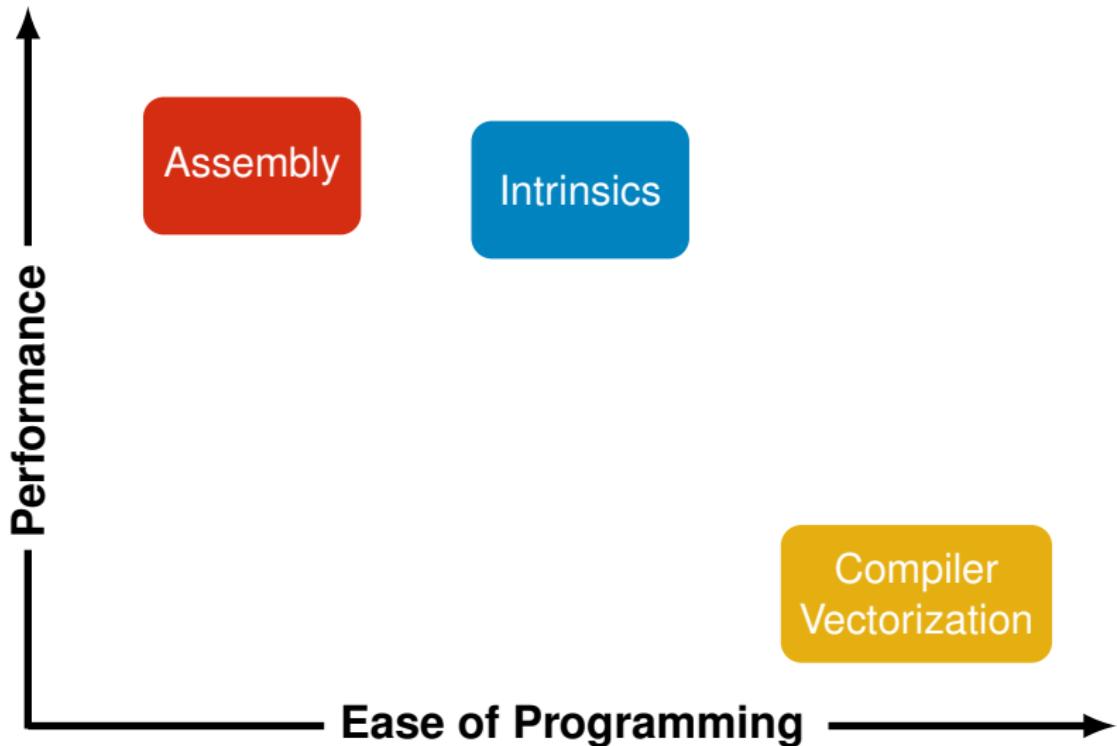


# Many- and Multi-core



- $\# \text{cores} \geq 52$
  - low frequency
  - weak/no Out-of-Order
  - better energy efficiency
  - high Flop/s
- 
- $\# \text{cores} \leq 18$
  - high frequency
  - deep Out-of-Order
  - large shared cache

# Vectorization techniques



## HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

## HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

complex 3-dim vector

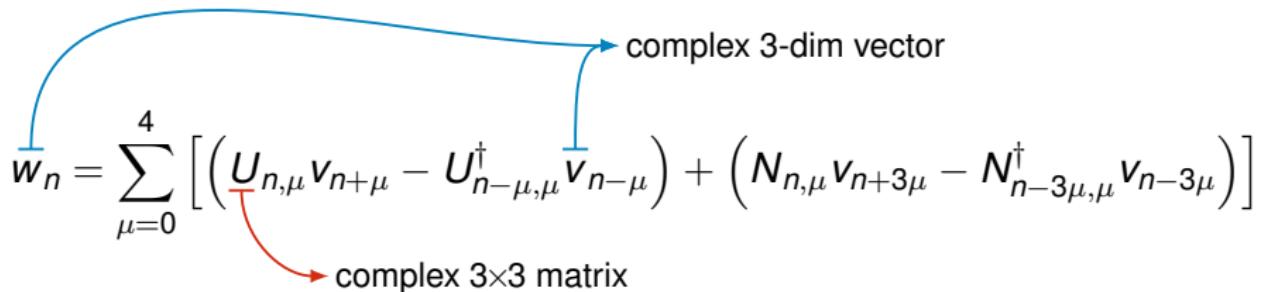
## HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

complex 3-dim vector

complex  $3 \times 3$  matrix



# HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

complex 3-dim vector

complex  $3 \times 3$  matrix

$U(3)$  matrix  
↪ reconstruct from 14 floats

# HISQ Dslash

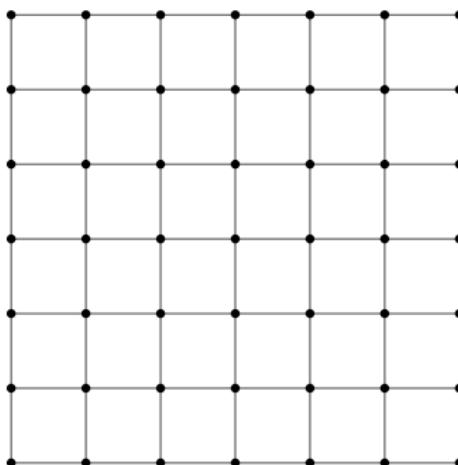
$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

↗ complex 3-dim vector  
↗ complex  $3 \times 3$  matrix  
↗  $U(3)$  matrix  
↪ reconstruct from 14 floats

$\mu$   
↑  
 $\nu$

↑ matrix  
● vector



# HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

complex 3-dim vector

complex  $3 \times 3$  matrix

U(3) matrix  
↪ reconstruct from 14 floats

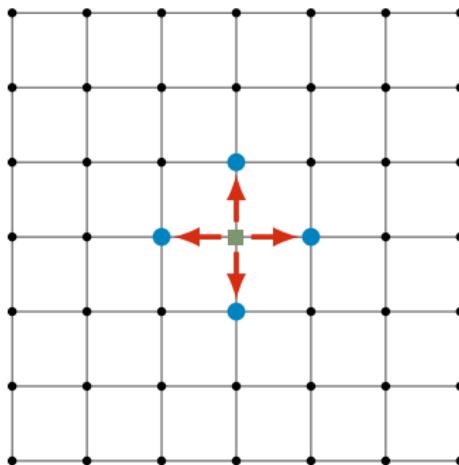
$\mu$

$\nu$

matrix

vector

$w_\square = \text{standard}$



# HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

complex 3-dim vector

complex  $3 \times 3$  matrix

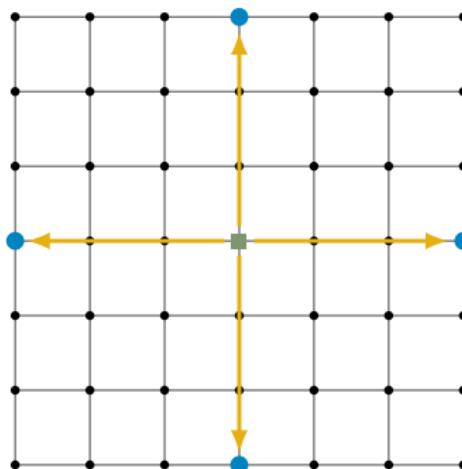
U(3) matrix  
↪ reconstruct from 14 floats

$\mu$   
↑  
 $\nu$

matrix  
● vector

$w_\square =$  standard  
+ naik

↪ precalculated  
three-link term



# HISQ Dslash

$$w = Dv$$

$$w_n = \sum_{\mu=0}^4 \left[ \left( U_{n,\mu} v_{n+\mu} - U_{n-\mu,\mu}^\dagger v_{n-\mu} \right) + \left( N_{n,\mu} v_{n+3\mu} - N_{n-3\mu,\mu}^\dagger v_{n-3\mu} \right) \right]$$

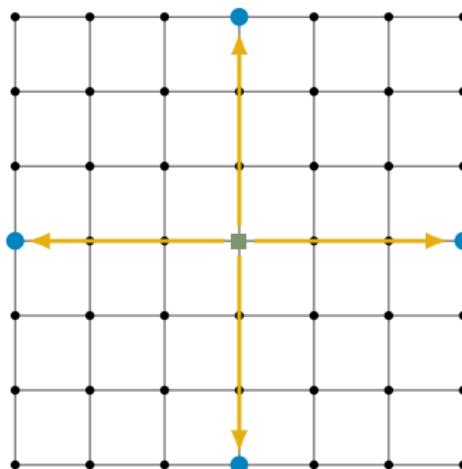
complex 3-dim vector  
complex  $3 \times 3$  matrix  
 $U(3)$  matrix  
 $\hookrightarrow$  reconstruct from 14 floats

$\mu$   
 $\nu$

↑ matrix  
● vector

$w_\square =$  standard  
+ naik

$\hookrightarrow$  precalculated three-link term



1146 Flop/site

0.8 Flop/byte  
 $\hookrightarrow$  single-precision

# Multiple right-hand sides

---

---

Memory

---

Memory

---

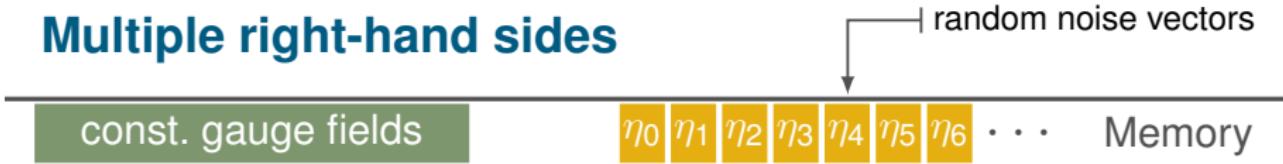
# Multiple right-hand sides

const. gauge fields

Memory

Memory

## Multiple right-hand sides



## Multiple right-hand sides



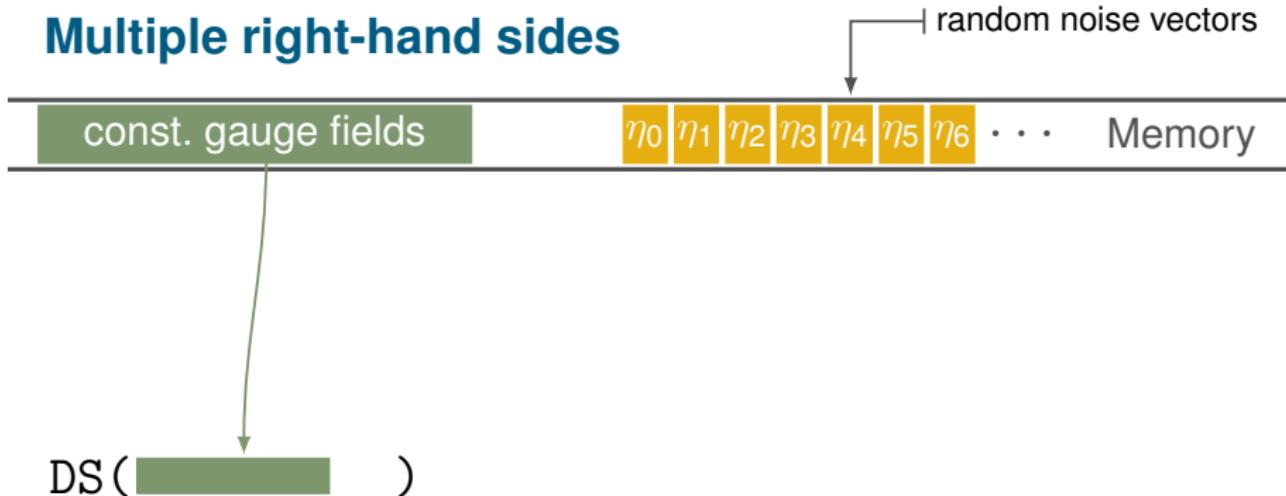
DS( )

---

Memory

---

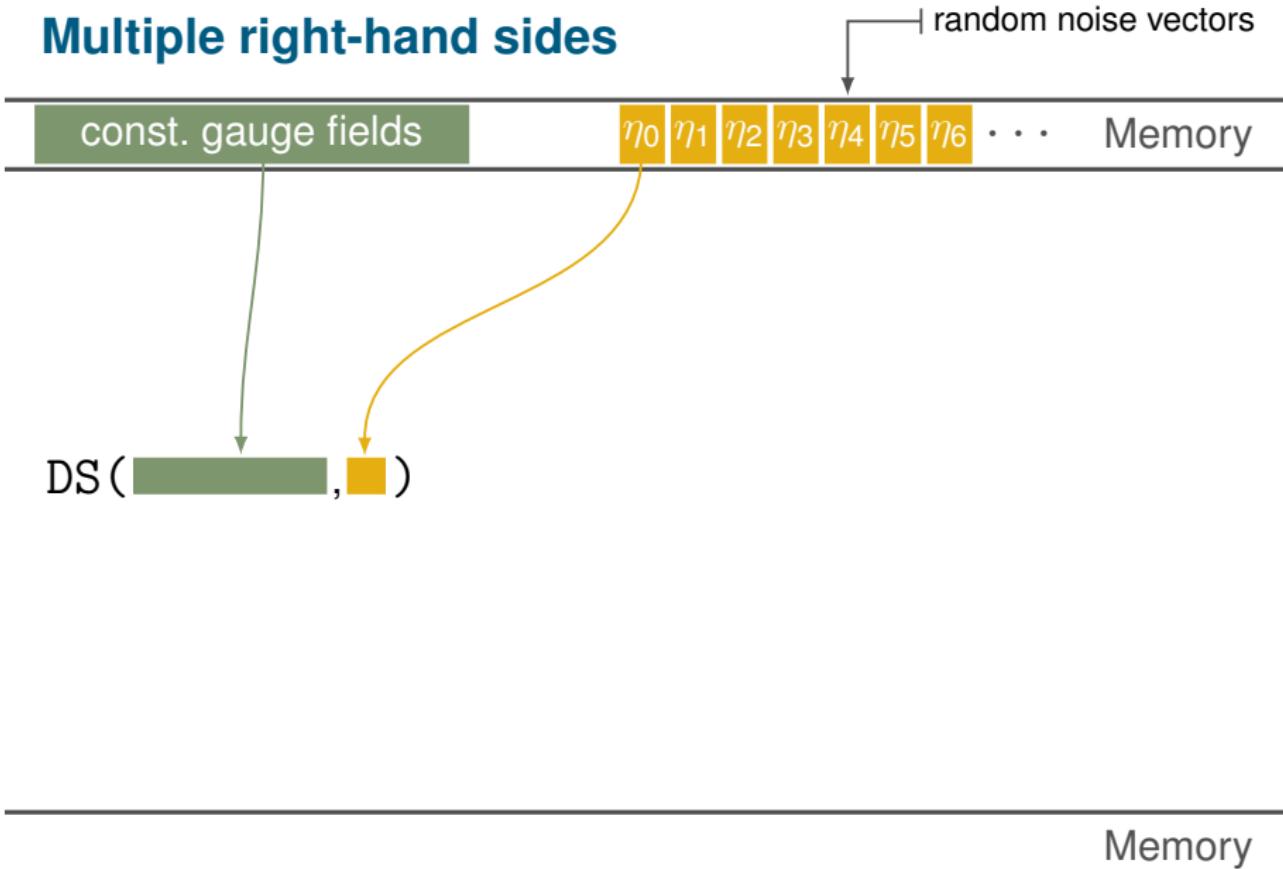
## Multiple right-hand sides



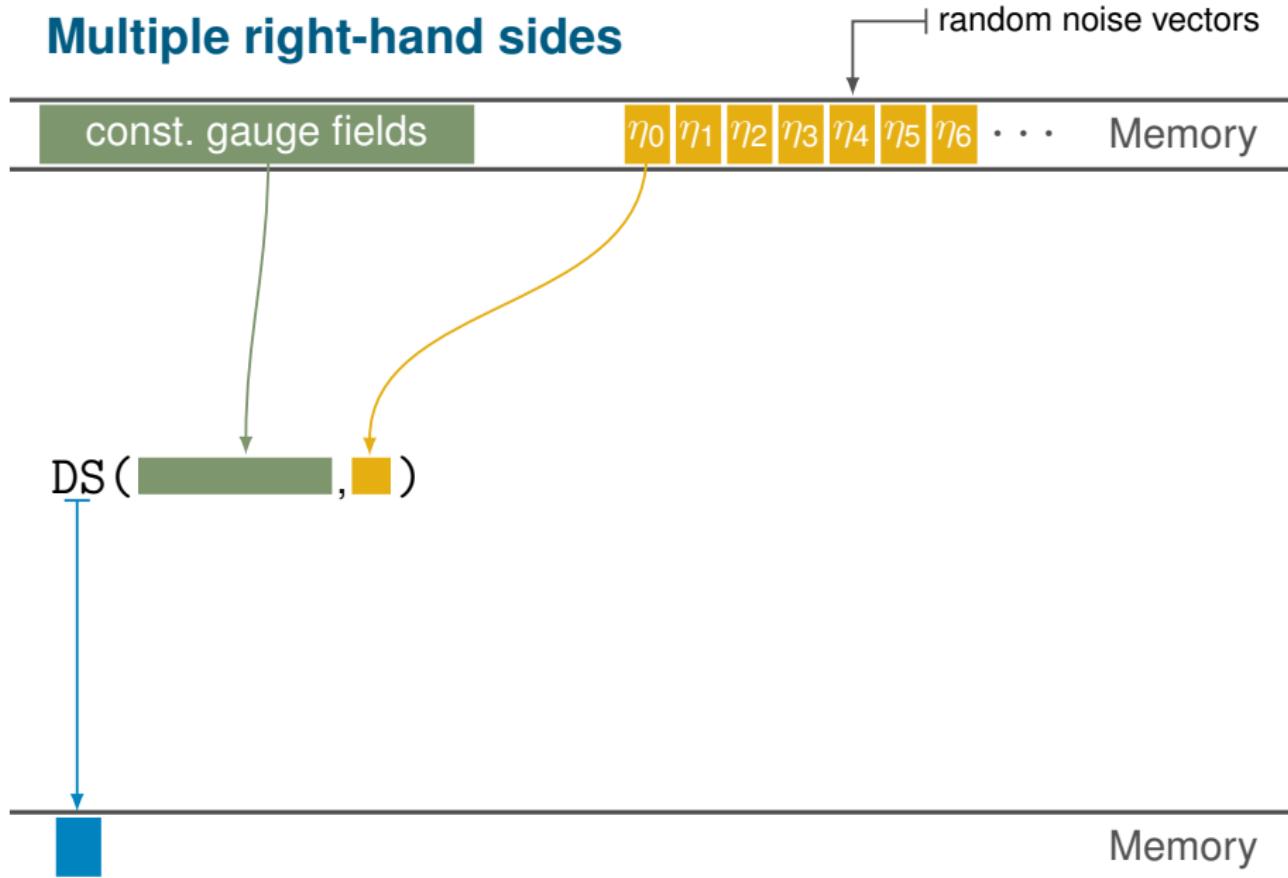
---

Memory

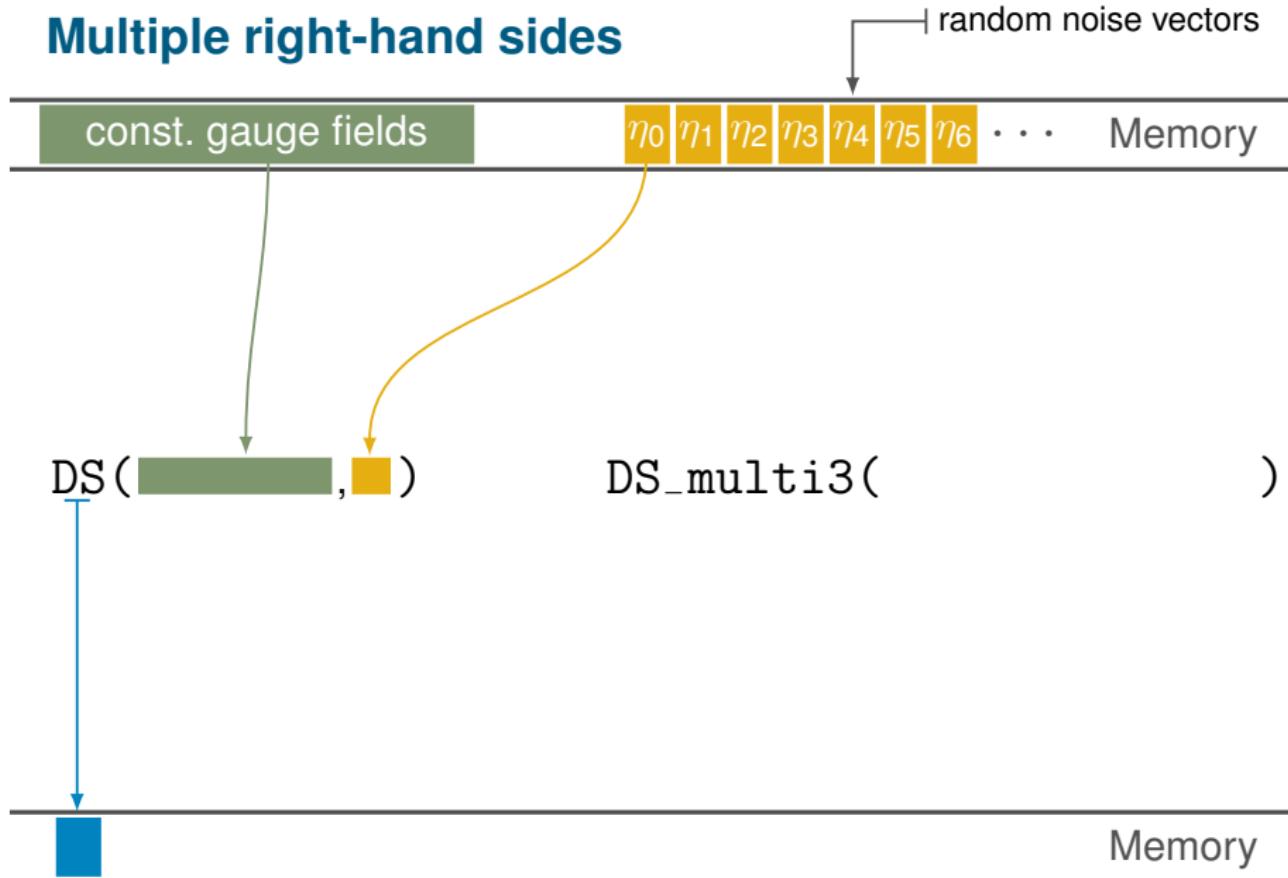
## Multiple right-hand sides



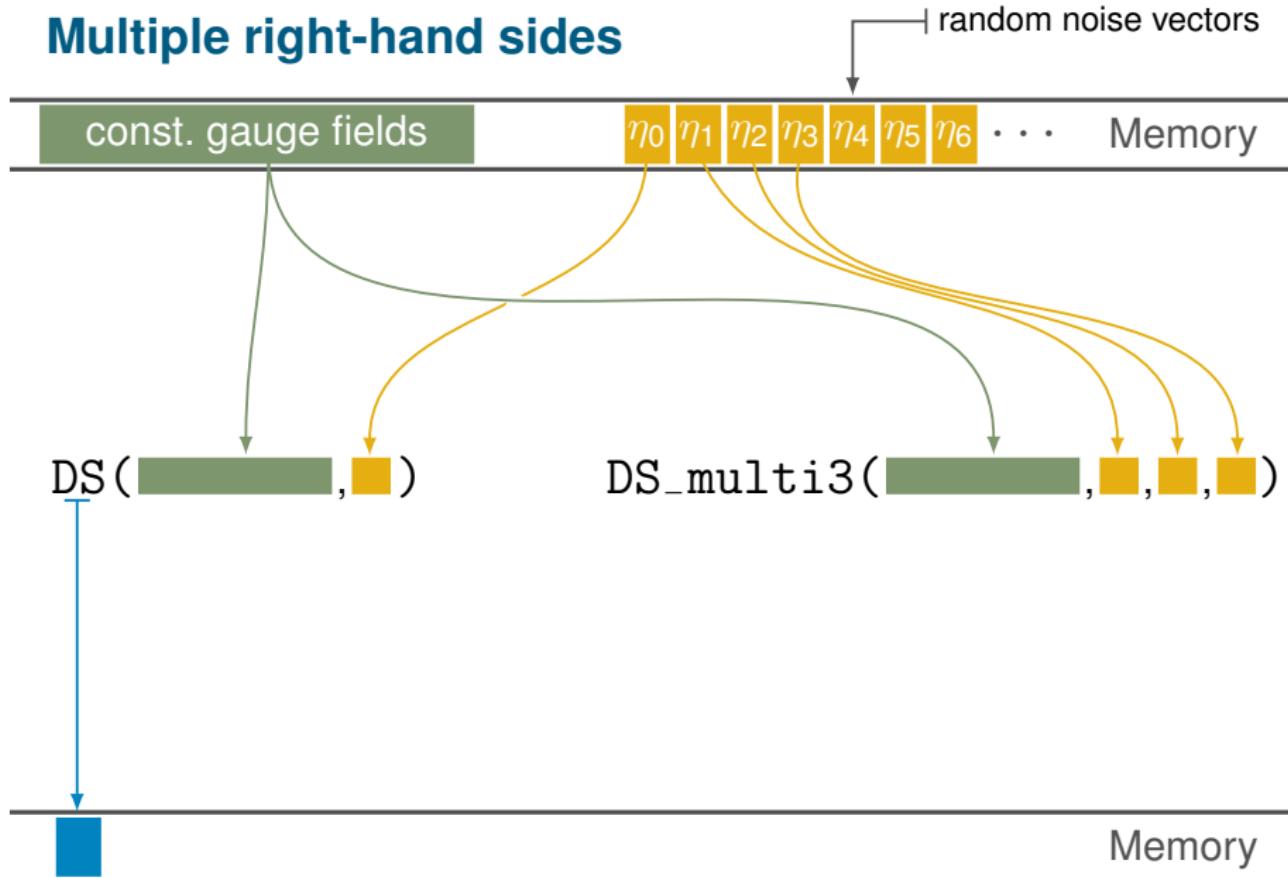
## Multiple right-hand sides



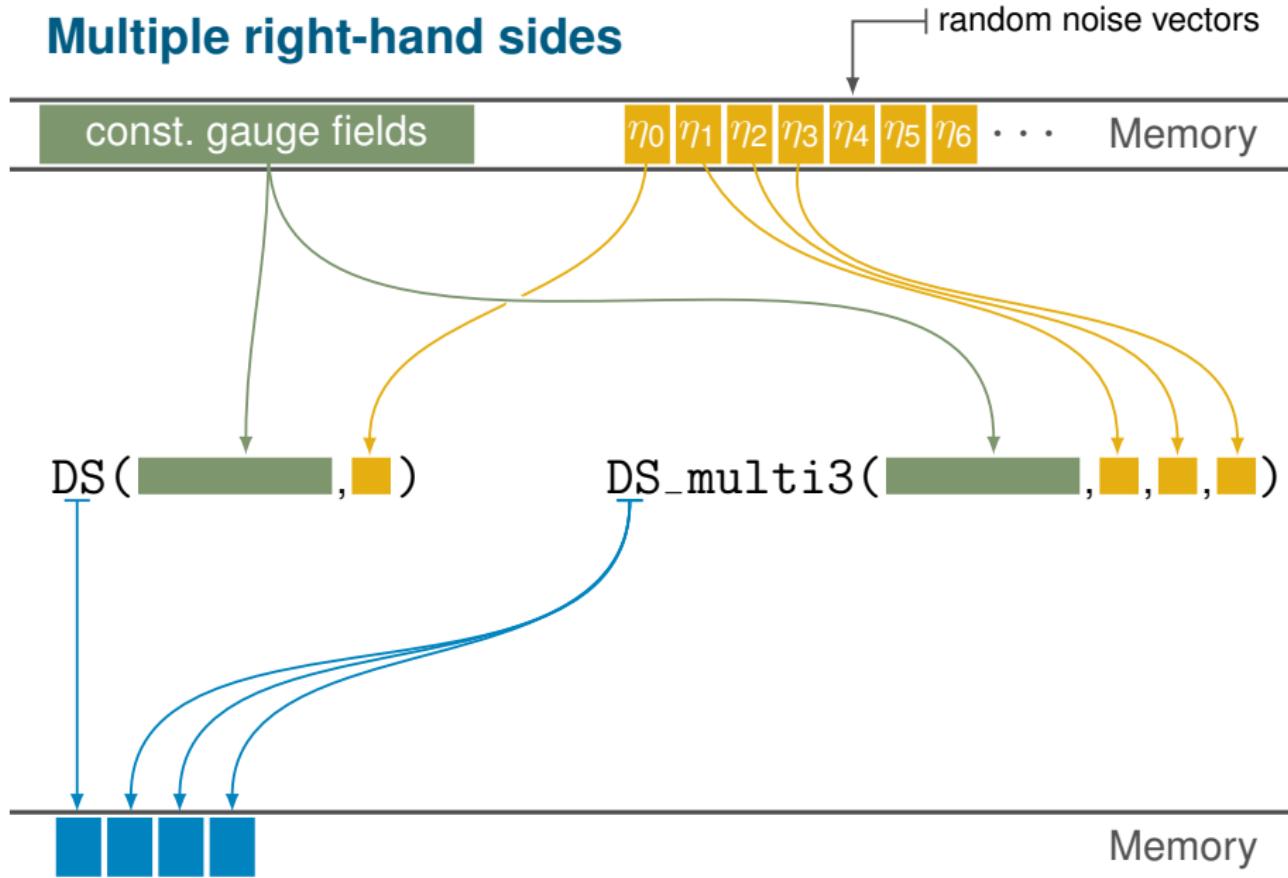
## Multiple right-hand sides



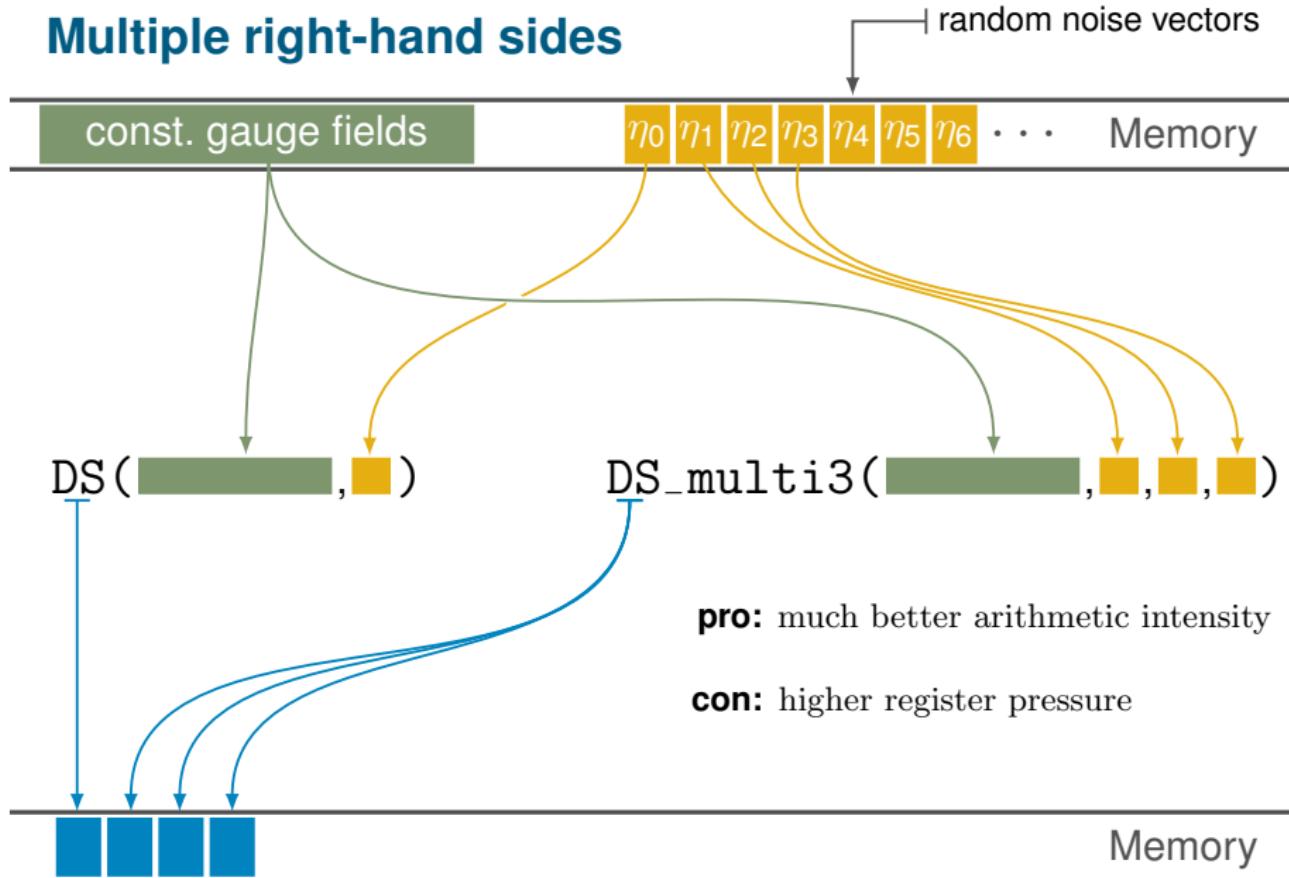
## Multiple right-hand sides



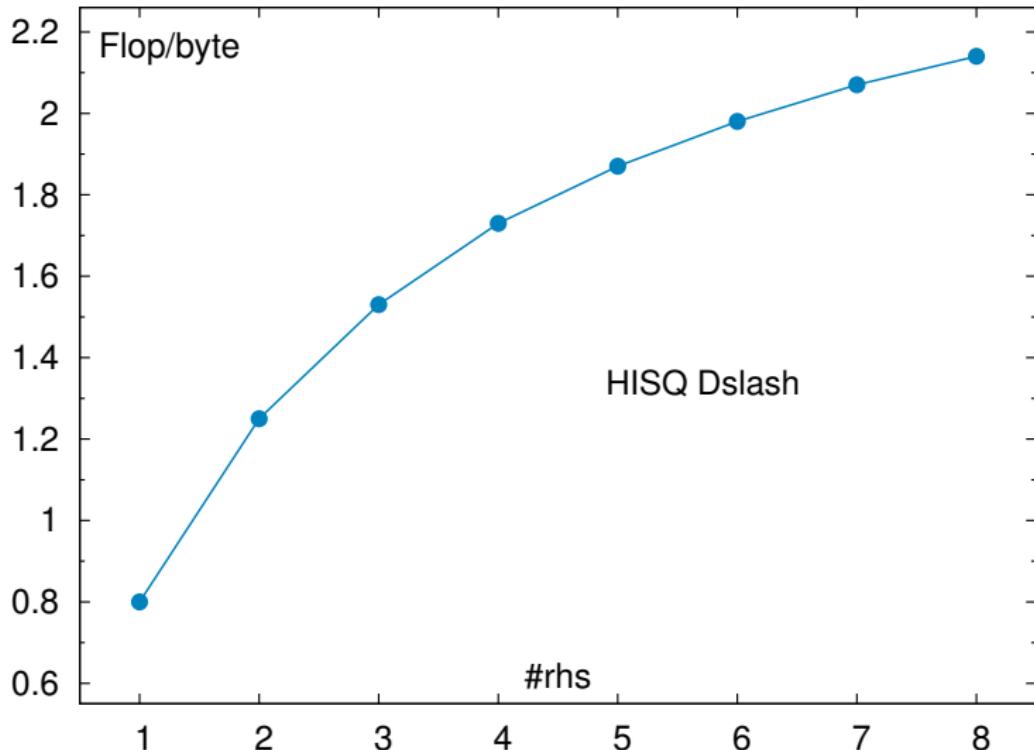
## Multiple right-hand sides



## Multiple right-hand sides

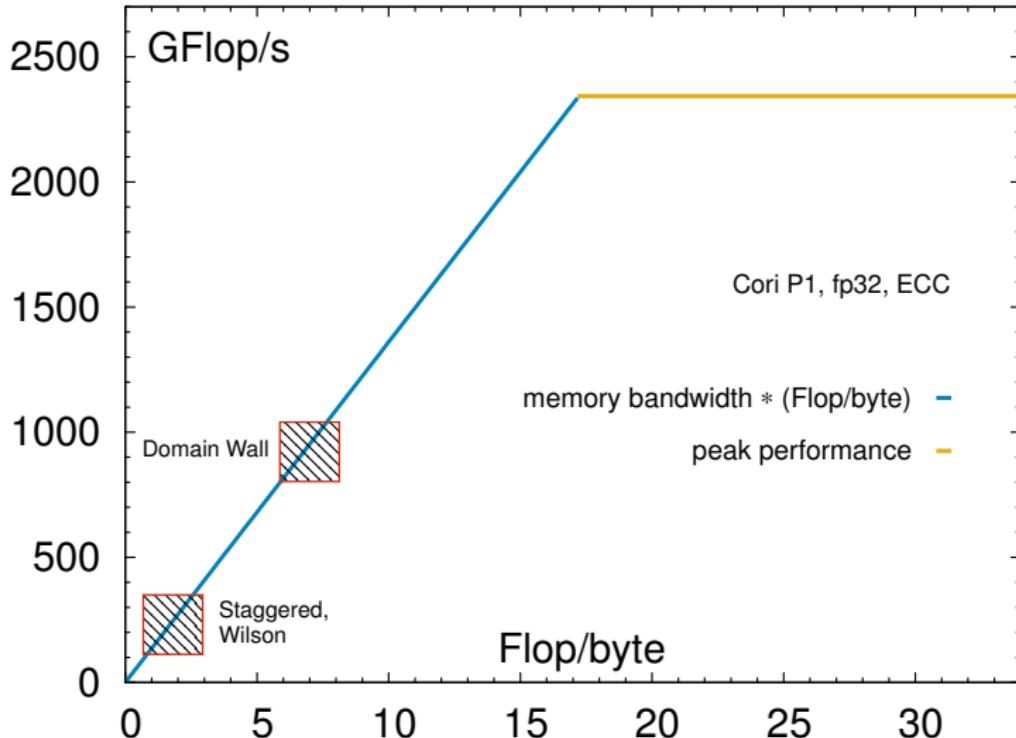


## Multiple right-hand sides (rhs)



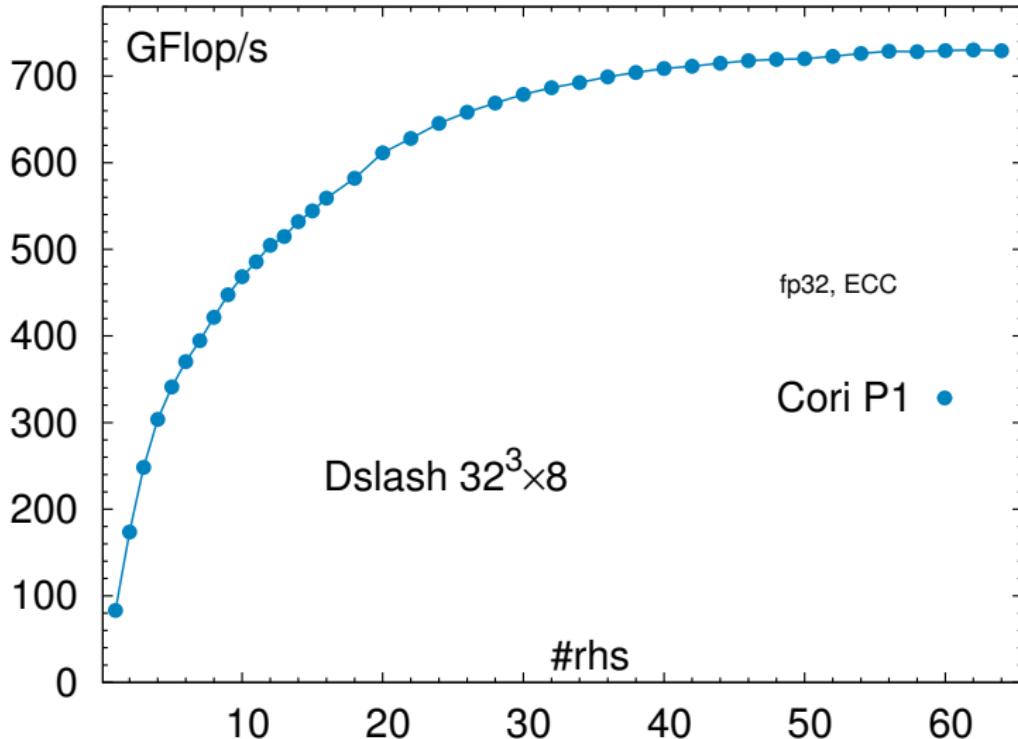
# Performance on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)



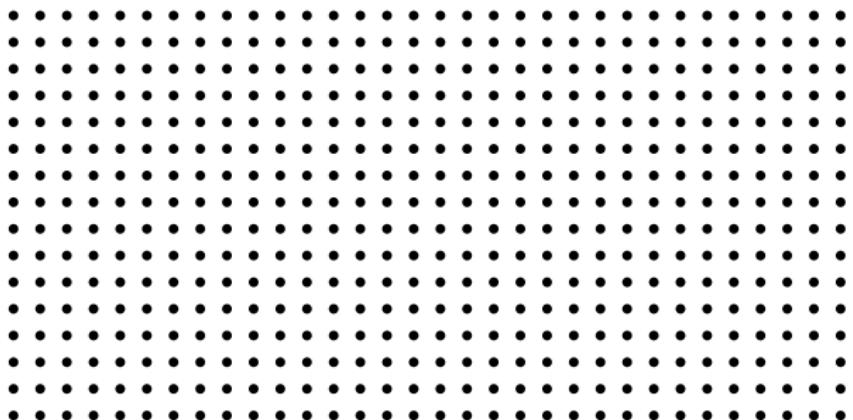
# Dslash on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)



# Vectorizing QCD

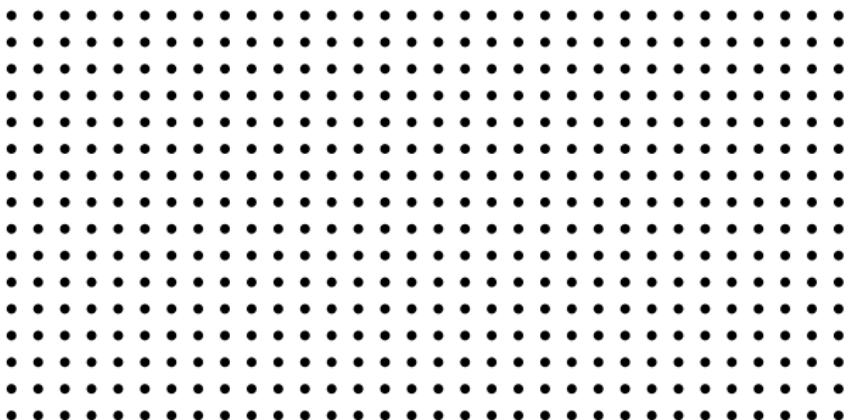
## Site fusion



# Vectorizing QCD

## Site fusion

- even
- odd

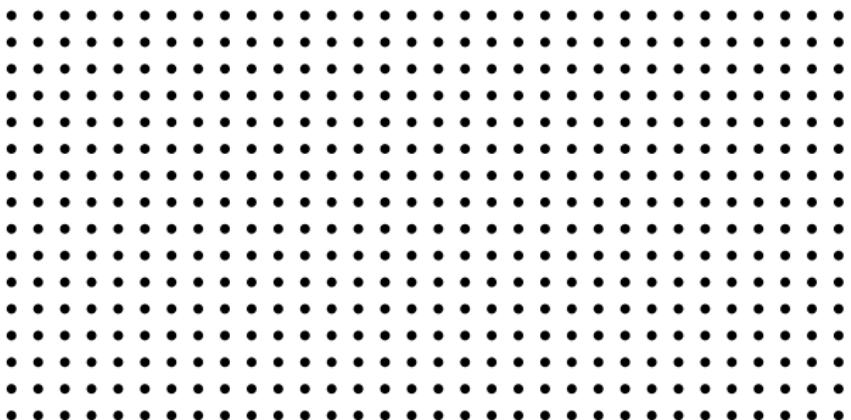


- even-odd preconditioning
  - need to fuse sites of the same parity

# Vectorizing QCD

## Site fusion

- even
- odd

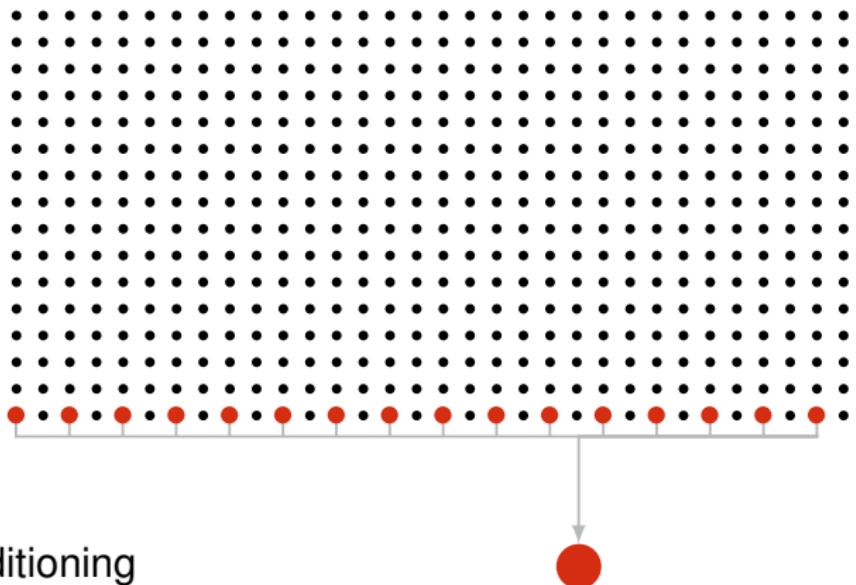


- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse e.g. 16 sites (fp32, AVX512)

# Vectorizing QCD

## Site fusion

- even
- odd

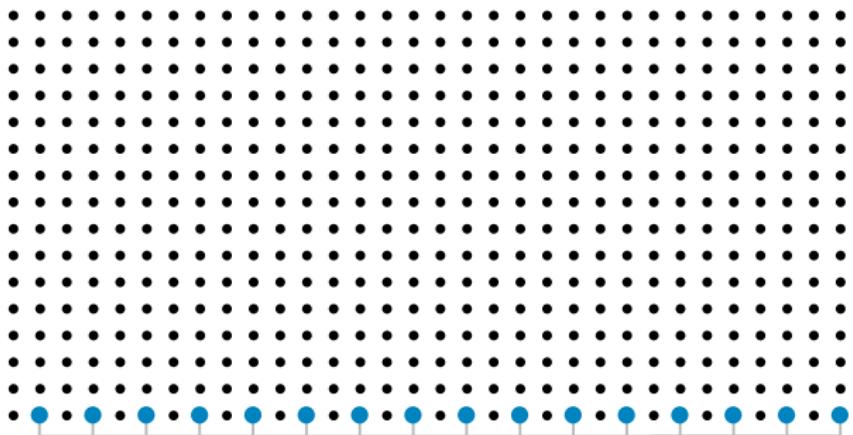


- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse e.g. 16 sites (fp32, AVX512)

# Vectorizing QCD

## Site fusion

- even
- odd

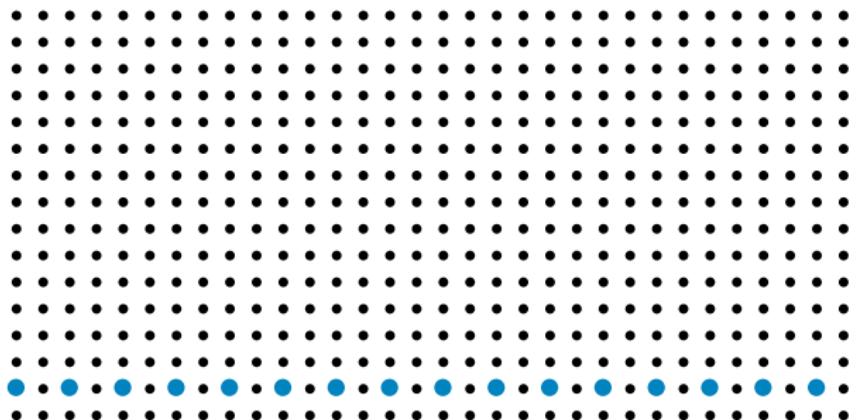


- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse e.g. 16 sites (fp32, AVX512)

# Vectorizing QCD

## Site fusion

- even
- odd



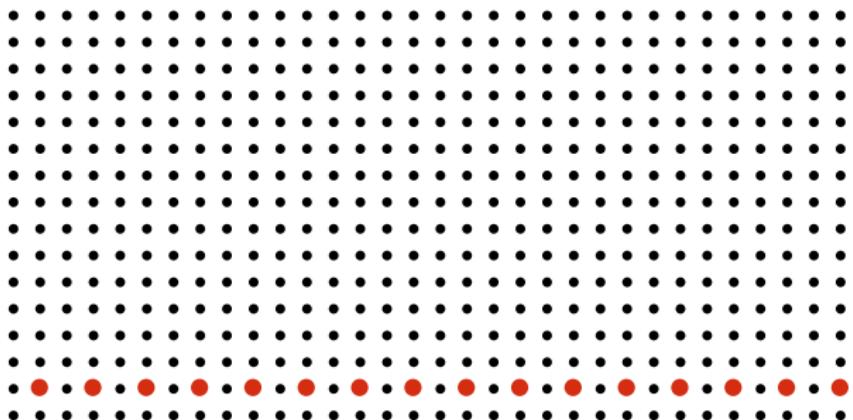
- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse e.g. 16 sites (fp32, AVX512)



# Vectorizing QCD

## Site fusion

- even
- odd



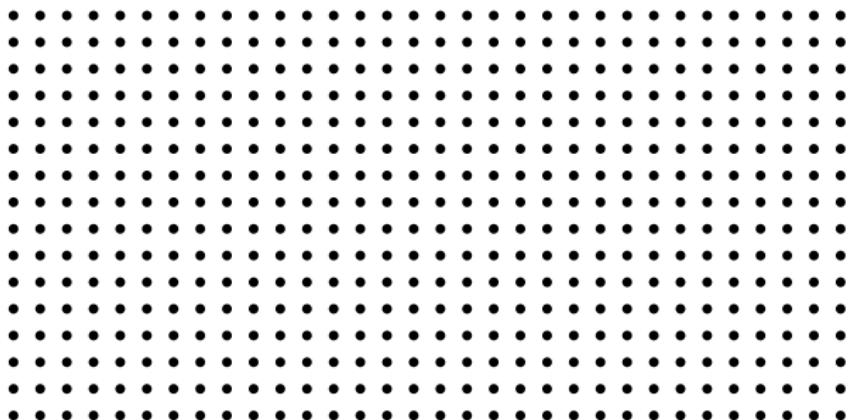
- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse e.g. 16 sites (fp32, AVX512)



# Vectorizing QCD

## Site fusion

- even
- odd

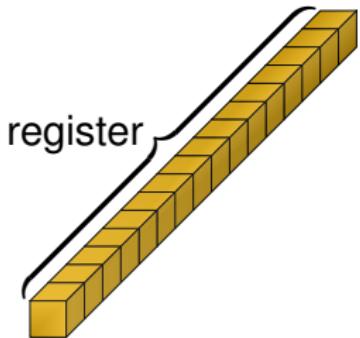


- even-odd preconditioning
  - need to fuse sites of the same parity
- fuse e.g. 16 sites (fp32, AVX512)
  - $32^3 \times 8 \longrightarrow 2 \times 32^2 \times 8$
  - constraints on lattice size



# Vectorizing QCD

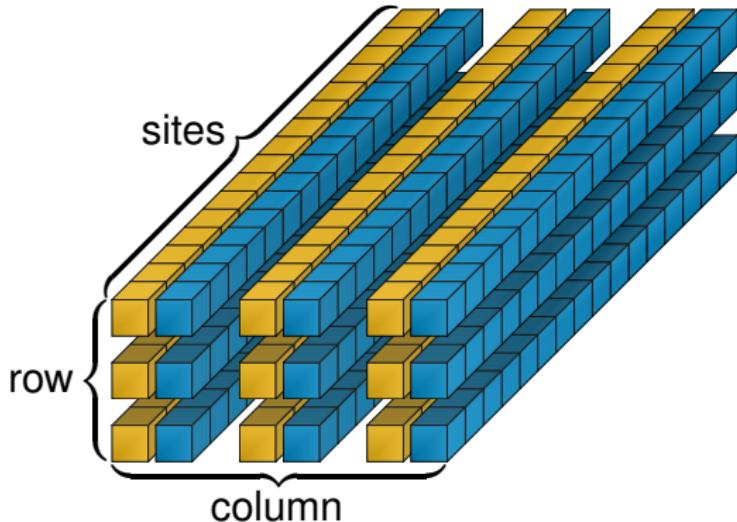
## Matrix-vector multiplication



# Vectorizing QCD

## Matrix-vector multiplication

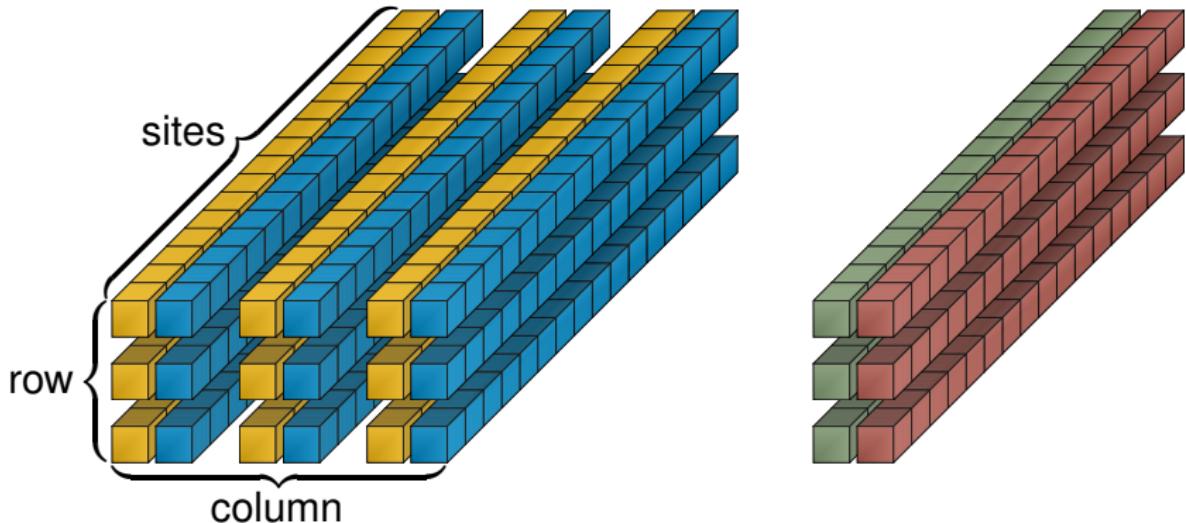
real      imag  
matrix    █    █



# Vectorizing QCD

## Matrix-vector multiplication

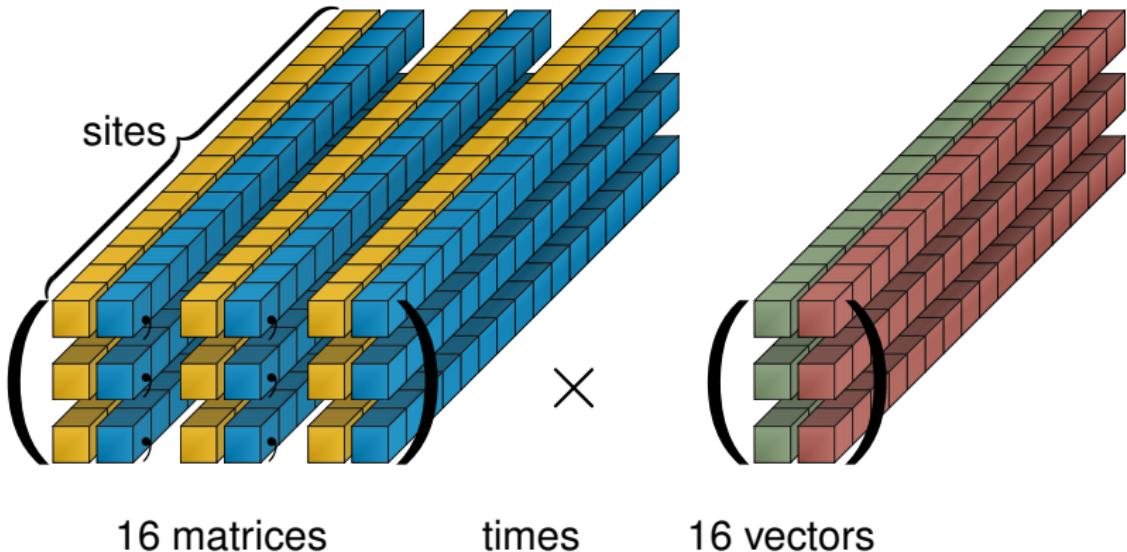
	real	imag
matrix	yellow	blue
vector	green	red



# Vectorizing QCD

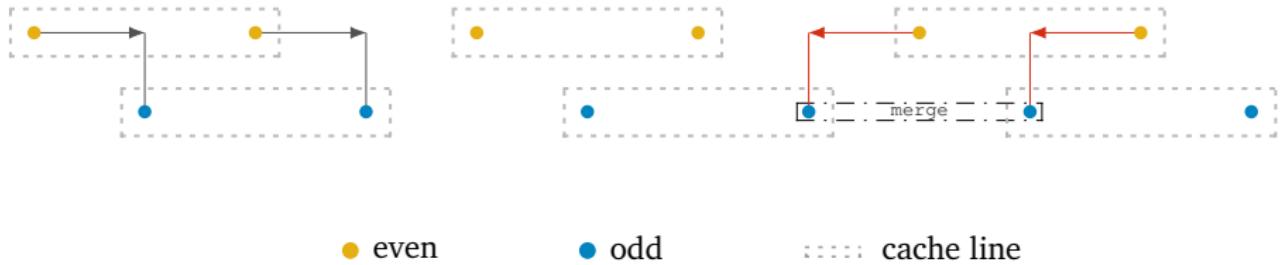
## Matrix-vector multiplication

	real	imag
matrix	Yellow	Blue
vector	Green	Red

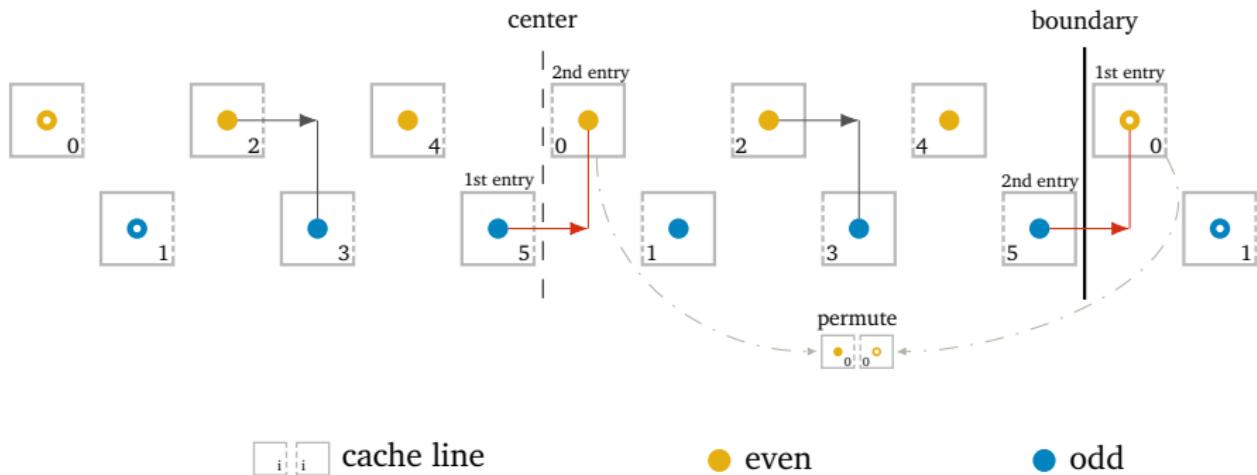


- very natural implementation → use SIMD registers like scalars

# Naive site fusion



# Efficient site fusion



# Performance gains on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)

Dslash,  $32^3 \times 8$ , 16 right-hand sides

optimizations

speedup

1 core, scalar



1x

# Performance gains on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)

Dslash,  $32^3 \times 8$ , 16 right-hand sides

optimizations

speedup



# Performance gains on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)

Dslash,  $32^3 \times 8$ , 16 right-hand sides

optimizations

speedup



# Performance gains on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)

Dslash,  $32^3 \times 8$ , 16 right-hand sides

## optimizations

## speedup

1 core, scalar		1x
32 cores, scalar		27x
1 core, vectorized		25x

# Performance gains on Haswell

(dual-socket E5-2698V3, 2.3 GHz, 32 cores)

Dslash,  $32^3 \times 8$ , 16 right-hand sides



Thank you for your attention!

# C

```
void add( double *a, double *b,
          double *c ) {

    for ( int i = 0; i < 8; ++i ) {

        c[i] = a[i] + b[i];
    }
}
```

# Intrinsics

```
void add( double *a, double *b,
          double *c ) {

    __m512d a_, b_, c_;

    a_ = _mm512_load_pd( a );
    b_ = _mm512_load_pd( b );

    c_ = _mm512_add_pd( a_, b_ );
    _mm512_store_pd( c, c_ );
}
```

# Assembly

```
void add( double *a, double *b,
          double *c ) {

    __asm {

        mov rdi, a
        mov rsi, b
        mov rdx, c

        vmovapd zmm1, [rdi]
        vmovapd zmm2, [rsi]
        vaddpd zmm3, zmm2, zmm1
        vmovapd [rdx], zmm3
    }
}
```